Red Hat Enterprise Virtualization 3.5 Technical Notes

Technical Notes for Red Hat Enterprise Virtualization 3.5 and Associated Packages

Red Hat Enterprise Virtualization Documentation Team

Red Hat Enterprise Virtualization 3.5 Technical Notes

Technical Notes for Red Hat Enterprise Virtualization 3.5 and Associated Packages

Red Hat Enterprise Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@redhat.com

Legal Notice

Copyright © 2015 Red Hat, Inc..

This document is licensed by Red Hat under the Creative Commons Attribution-ShareAlike 3.0 Unported License. If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

These Technical Notes provide documentation of the changes made between release 3.4 and release 3.5 of Red Hat Enterprise Virtualization. Subsequent advisories that provide enhancements, provide bug fixes, or address security flaws are also listed. They are intended to supplement the information contained in the text of the relevant errata advisories available via Red Hat Network.

Table of Contents

Introduction 2

Chapter 1. RHSA-2015:0888 Moderate: Red Hat Enterprise Virtualization Manager 3.5.1 3

Chapter 2. RHBA-2015:0161 ovirt-hosted-engine-setup 8

Chapter 3. RHEA-2015:0160 ovirt-node 11

Chapter 4. RHBA-2015:0159 vdsm 16

Chapter 5. RHSA-2015:0158 Moderate: Red Hat Enterprise Virtualization Manager 3.5.0 ... 20

Revision History 38

Introduction

These Technical Notes provide documentation of the changes made between release 3.4 and release 3.5 of Red Hat Enterprise Virtualization. They are intended to supplement the information contained in the text of the relevant errata advisories available via Red Hat Network. Red Hat Enterprise Virtualization 3.x errata advisories are available at https://rhn.redhat.com/errata/rhel6-rhev-errata.html.

A more concise summary of the features added in Red Hat Enterprise Virtualization 3.5 is available in the Red Hat Enterprise Virtualization 3.5 Manager Release Notes.

Chapter 1. RHSA-2015:0888 Moderate: Red Hat Enterprise Virtualization Manager 3.5.1

The bugs contained in this chapter are addressed by advisory RHSA-2015:0888. Further information about this advisory is available at https://rhn.redhat.com/errata/RHSA-2015:0888.html.

ovirt-engine-backend

BZ#1176552

Previously, when a user attached a storage domain that was already in use and managed by a different Manager, no warning was provided, and the action lead to potential metadata corruption. With this update, users are notified with a warning message that the intended storage domain is already attached to another data center managed in another Red Hat Enterprise Virtualization environment. The user can choose to continue and overwrite the metadata or to cancel the action.

BZ#1184807

Previously, less than or equal to (<=) was used for storage thresholds. In addition, integer numbers were used and caused fractions to be truncated. This triggered alerts for low disk space when it shouldn't have. With this update, when checking thresholds, less than (<) is now used with decimal points taken into account. Alerts for low disk space are now generated appropriately.

BZ#1192014

When installing a new host using the Virt mode (without Gluster support), port 111 was not opened in TCP and UDP and blocked rpc.statd. With this update, the required ports are opened in the firewall.

BZ#1176546

Previously, because OVF data of virtual machines or templates with no disks was not stored on any OVF store, when you detach a storage domain and attach the domain to another data center, these virtual machine or templates got lost. With this update, OVF data of diskless virtual machines or templates is stored on all OVF stores on all domains.

BZ#1195000

Previously, removing a snapshot that contains only a memory volume (i.e. live snapshot without disks) left the snapshot locked in the database. With this update, such snapshots are removed successfully.

BZ#1178646

Previously, when trying to attach an imported storage domain that carried existing metadata to an uninitialized data center, an exception from VDSM was returned with no proper warning message. With this update, the storage domain is checked for existing metadata and an error message is provided to advise users to attach a clean storage domain first.

ovirt-engine-cli

BZ#1181681

This update adds the call 'isattached' to the REST API for detecting whether a storage domain is attached to a data center before attempting to import the storage domain. This functionality allows users to check if a storage domain is already attached to a storage pool before importing it to a new environment, thereby preventing corruption in the data of a storage domain that is already activated in a different environment.

IMPORTANT: When executed, the call causes the storage domain to become disconnected from the host where the call is executed.

ovirt-engine-dwh

BZ#1181642

Previously, if connection to the engine database failed or was lost temporarily, the job that checks the 'disconnectDwh' flag did not restore the connection. As a result, the ETL service, ovirt-engine-dwhd, remained running. Now, a process has been added that supports the attempt to reconnect to the database, and the 'disconnectDwh' flag is checked correctly.

BZ#1180867

A problem with starting the oVirt-ETL (extract, transform, load) service caused installation of the data warehouse to fail. Now, the data warehouse uses the engine's script /usr/share/ovirt-engine/bin/java-home to detect the JAVA_HOME location, so this error no longer occurs.

ovirt-engine-iso-uploader

BZ#1188326

Previously, engine-iso-uploader was hardcoded to connect to the ISO storage domain as 'localhost', which prevented remote uploading of ISO images from another host via SSH. Now, engine-iso-uploader properly retrieves the hostname for the local storage domain with a REST API request and it is possible to remotely upload an ISO image to a ISO storage domain via SSH.

ovirt-engine-log-collector

ovirt-engine-log-collector

BZ#1175137

SOS 3 uses a difference plugin scheme from SOS 2 to account for differences between Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7. This resulted in some information not being collected on hosts. Now, additional plugins have been configured so that the required information is collected.

ovirt-engine-sdk

BZ#1182158

This enhancement adds the ability to import block storage domains using the REST API.

ovirt-engine-setup

BZ#1188971

The Manager's Java Virtual Machine heap size was configured to be 1GB by default. Large setups made the Manager run out of heap memory and required manual configuration to increase the heap size. With this update, engine-setup now automatically configures the heap size to be either a minimum of 1GB or 1/4 of the available memory. Large setups only need to use a machine with enough memory (the recommendation is 16 GB of ram), and heap size configuration will be done automatically, thus preventing out of memory failures.

BZ#1192954

Previously, lc_messages might have been set to a non-English locale in postgresql.conf. The engine-backup --mode=restore command did not filter expected errors, which were in English, and failed. With this update, engine-backup --mode=restore was changed to require lc_messages to be 'en_US.UTF-8'. As a result, if lc_messages is not 'en_US.UTF-8', a more helpful error message is returned.

BZ#1196136

With this update, zombie commands are cleaned to avoid getting stuck waiting for task and commands completion.

BZ#1197616

Previously, when creating a template and then running an upgrade command, an implementation that was supposed to only clear zombie tasks and commands also killed pending tasks for template creation. This caused the template creation to be stuck. With this update, only zombie tasks and commands are cleared as normal.

ovirt-engine-userportal

BZ#1194394

Previously, certain "form-filler" single sign-on (SSO) solutions (such as Indeed ID) that did not fire "change" events would not work with the Red Hat Enterprise Virtualization login pages. "Form-filler" SSO solutions that did fire "change" events (such as LastPass) worked correctly. With this update, the login forms have been enhanced to work without "change" events. All form-filler SSOs now work.

ovirt-hosted-engine-setup

BZ#1190636

Previously, hosted engine used the vdsClient utility to communicate with VDSM, which meant that SSL would timeout on sync commands that would take more than 60 seconds to complete. Now, the vdscli library is used for storage operations, due to configurable timeout, and longer sync commands no longer fail due to SSL timeout.

BZ#1192462

Previously, hosts were hardcoded to overwrite the iptables rules when the host was added using the 'hosted-engine --deploy' command, even if the user answered 'No' to the question 'iptables was detected on your computer, do you wish setup to configure it?'. Now, the host is not hardcoded and an answer of 'No' to this question is recognised by both the 'hosted-engine --deploy' configuration as well as during the request to the engine to add the host. Therefore, answering 'No' prevents the existing iptables rules from being overwritten on the host.

BZ#1181585

Previously, the hosted engine would check to ensure an ISO image was readable by the VDSM user but not necessarily the KVM user. This meant it was possible for the virtual machine to be unable to boot from the ISO even if though it passed the check. Now, this check has been expanded and, if the ISO image passes, the virtual machine can boot from it as expected

RFEs

BZ#1196199

With this feature, you can now delete virtual machine disk snapshots from running virtual machines.

BZ#1194272

Previously, administrators can not define explicit permissions for users on who can use live storage migration. With this update, a new permission called DISK_LIVE_STORAGE_MIGRATION which will be part of DISK_STORAGE_MANIPULATION was introduced so administrators can manipulate which users will be able to use live storage migration and which can not.

BZ#1187985

This feature add default options for Drac7 fencing agent.

BZ#1186375

Previously, it was only possible to import pre-exisiting export domains type of NFS type. With this update, it is now possible to import pre-existing Gluster and POSIX filesystem export domains as well.

BZ#1174814

Previously, all Windows versions used 'sysprep.inf' as the file name for sysprep files but since newer Windows versions uses 'Unattend.xml' as the file name, the file names were mismatched for newer Windows versions. This RFE sets the correct file name for sysprep files of newer Windows versions.

vulnerability

BZ#1189085

It was discovered that a directory shared between the ovirt-engine-dwhd service and a plug-in used during the service's startup had incorrect permissions. A local user could use this flaw to access files in this directory, which could potentially contain sensitive information.

Chapter 2. RHBA-2015:0161 ovirt-hosted-engine-setup

The bugs contained in this chapter are addressed by advisory RHBA-2015:0161. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2015-0161.html.

ovirt-hosted-engine-setup

BZ#1134873

Previously, hosted engine high availability services could run on the host even if they were unconfigured, which would prevent hosted engine deployment. Now, these services are stopped if they are running unconfigured and they do not interfere with hosted engine deployment.

BZ#1078206

Previously, during the hosted engine deployment, selecting a bonded interface to be used as the base for the 'rhevm' bridge would cause the deployment to fail. Now, it is possible to select a bonded interface for the 'rhevm' bridge during hosted engine deployment.

BZ#1043906

Previously, the hosted engine deployment would add additional hosts to the 'Default' cluster and hosts could not be added to the environment if this cluster had been renamed. Now, if the hosted engine deployment establishes that there is no cluster named 'Default' the user is prompted for the cluster name to which the host will be added.

BZ#1125812

Deployment of the hosted engine is now supported on Red Hat Enterprise Linux 7 hosts.

BZ#1116785

The use of multicast MAC addresses is not supported by libvirt. Previously, the hosted engine deployment wizard was not verifying the validity of a provided MAC address and the wizard was failing with an unexpected error. Now, MAC addresses are verified and the user is prompted to enter the address again if it is not valid.

BZ#1109929

Previously, if the host could not be added to the Default cluster during hosted engine deployment, the deployment would fail with an ambiguous error. Now, the issue is properly detected and a specific error is raised to make the user aware of the problem.

BZ#1168267

Previously, failure to connect to the engine API for any reason during hosted engine deployment would report incorrectly that the host could not be added to the cluster. Now, the reporting of these exceptions has been improved and the user is provided a more coherent report of the failure.

BZ#1103672

Previously, if the NX flag was not checked in the BIOS of certain Intel CPU types, which require NX as well as VMX to support virtualization, the deployment script for the hosted engine would not accurately detect the CPU type and the deployment would stall. Now, the user is prompted to check the NX flag in the system BIOS if the CPU type is not accurately detected, and the deployment exits gracefully.

BZ#1076944

Previously, during the hosted engine deployment, selecting a VLAN-tagged network interface to be used as the base for the 'rhevm' bridge would cause the deployment to fail. Now, it is possible to select a VLAN-tagged network interface for the 'rhevm' bridge during hosted engine deployment.

BZ#1106556

Previously, if the deployment of the hosted engine was aborted after the engine virtual machine had been created, a subsequent deployment would fail until the engine virtual machine had been manually destroyed. Now, an option has been added so that the engine virtual machine is destroyed when the deployment is aborted.

BZ#1105249

Deployment of the hosted engine via ssh without terminal mode fails when establishing the storage connection. Attempting to deploy the hosted engine in this way now fails with a warning that it should be executed with terminal mode.

BZ#1107772

Previously, the confirmation of installation settings in the hosted engine deployment wizard defaulted to 'No', which would require the whole deployment to be restarted if selected by accident. Now, the confirmation of installation settings defaults to 'Yes' and an answer file is created so that if the user does select 'No' the answer file can be used for skipping some of the setup.

BZ#1105479

Previously, deployment of the hosted engine failed if the

'GATEWAY' was specified in network configuration files and 'BOOTPROTO' was set to 'none' due to use of deprecated 'addNetwork' API in network setup. Now, the python vdsm.cli module is used in network setup and deployment of the hosted engine does not fail in this situation.

BZ#1172545

Previously, a bug in the regular expression prevented correct fetching of the IP address for VLAN interfaces, which in turn prevented the deployment of the hosted engine. Now, the regular expression used to fetch the IP address for VLAN interfaces has been corrected and the hosted engine can be deployed using VLAN.

Chapter 3. RHEA-2015:0160 ovirt-node

The bugs contained in this chapter are addressed by advisory RHEA-2015:0160. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2015-0160.html.

ovirt-node

BZ#894258

Previously, when a Red Hat Enterprise Virtualization Hypervisor was registered to Subscription Asset Manager using a proxy, the proxy details were not set in the Hypervisor console. This issue has now been fixed, and the proxy details display correctly.

BZ#920171

Network bonds can now be automatically configured during installation using the 'bond_setup=' and 'bond=' kernel arguments for auto-installation of the Red Hat Enterprise Virtualization Hypervisor.

BZ#960833

In previous versions of the Red Hat Enterprise Virtualization Hypervisor, once the NFSv4 domain was set using the text user interface, the entry could not be removed as the relevant entry in the idmpad.conf was not being properly cleared. This has now been corrected so that users can remove the NFSv4 domain after it has been set.

BZ#966302

This feature enables the default console device to be set from within the Red Hat Enterprise Virtualization Hypervisor setup TUI.

BZ#1169865

Previously, the ISO size for the Red Hat Enterprise Virtualization Hypervisor was restricted to 256MB. Now, that size has been increased to 4.70GB, increasing the minimum disk size requirements for RHEV-H 7.0 to 10GB.

BZ#1162445

Previously, the Red Hat Enterprise Virtualization Hypervisor failed to interpret the major version of an image and attempts to reinstall the Hypervisor would fail with an error. This has now been corrected so that the Hypervisor accepts a major version mismatch and reinstalls as expected.

BZ#1155957

Previously, a multipath regression caused machines to fail to boot USB media created using DD. This has now been corrected and the Hypervisor can be installed or reinstalled from USB media.

BZ#1053505

Previously, reinstalling the Red Hat Enterprise Virtualization Hypervisor to a device with multipath devices resulted in a kernel panic on some hardware devices. Now, disks selected for wiping are checked to see if they are listed more than once and the Hypervisor is reinstalled as expected.

BZ#1062515

Confirmation is now required for storage layout in the TUI installation of the Red Hat Enterprise Virtualization Hypervisor to help prevent data loss in case of an incorrectly selected disk.

BZ#1067355

Previously, an incorrect call from the Red Hat Enterprise Virtualization Hypervisor TUI to Subscription Manager was preventing the Hypervisor from attaching to the Satellite 6 server. This call has now been rewritten and the TUI can be used to subscribe the Hypervisor to the Satellite server as expected.

BZ#1095140

Local configurations of kdump now work as expected in Red Hat Enterprise Virtualization Hypervisor 7.0

BZ#1084528

Previously, Red Hat Enterprise Virtualization Hypervisor 6.5 installations on nodes without usable disks would fail with an incorrect error message regarding keyboard. Now, the Hypervisor installer handles a lack of usable disks gracefully, displaying a message that there are no valid boot devices and disabling the 'Continue' button.

BZ#1078608

Previously, in the Red Hat Enterprise Virtualization Hypervisor 6.5 installation TUI, providing a path for in the 'Other Device' field to boot the Hypervisor when a different device path had previously been selected caused a ValueError exception to be thrown. Now, the custom device parser has been fixed and providing a second device path in the 'Other Device' field works as expected.

BZ#882846

Previously, a livecd image on a USB disk would be filtered out by the Red Hat Enterprise Virtualization Hypervisor TUI if the Hypervisor was booted from PXE. Now, the USB disk is visible as expected in the TUI if the Hypervisor has been booted from PXE.

BZ#1084276

Previously, the image-minimizer tool was missing in ovirt-node packages, causing edit-node to skip minimization and resulting in a larger-than-necessary ISO. The image-minimizer tool is now shipped in the ovirt-node-minimizer subpackage and ISOs generated by edit-node are minimized as expected.

BZ#1095028

Aborting the media integrity check during Red Hat Enterprise Virtualization Hypervisor 7.0 boot causes system halt and failure to boot. Remove the kernel argument rd.live.check from the kernel command line to prevent the media check on boot.

BZ#1039233

Previously, customizing the Red Hat Enterprise Virtualization Hypervisor 6.5 ISO and opening an rpm/srpm/file manifest in the plugin page of the setup TUI caused the TUI to crash due to non-existent manifest file. The plugin parser has now been fixed so that the TUI displays an error message and does not crash.

BZ#1156343

Behavior changes in Red Hat Enterprise Linux 7 requires 'boot=' argument to configure bonds, which was causing auto-installation of the Red Hat Enterprise Virtualization Hypervisor 7 to fail when user specified 'boot_setup=' argument, as in Hypervisor 6. These bonding changes are now supported in the Hypervisor and auto-installation of the Hypervisor succeeds as expected with 'bond={bond_name}:{list_of_interfaces}' argument.

BZ#1152948

A multipath regression meant multipath was incorrectly claiming devices, which prevented the Red Hat Enterprise Virtualization Hypervisor to boot because it could not locate the root file system. This has been adjusted so that multipath will only claim multipath devices and the Hypervisor boots as expected.

BZ#1063395

Previously, the Red Hat Enterprise Virtualization Manager would sometimes report a critical disk space error: "Critical, Low disk space. Host <host> has less than 500 MB of free space left on: /var/log" even when this was not the case. Now, the Manager reports disk space errors accurately.

BZ#1062111

Previously, the methodology for determining CPU Family used by the Red Hat Enterprise Virtualization Manager differed from that used by the Hypervisor, which yielded inconsistent results. Now, the CPU Family has been dropped and only the correct CPU Model name is displayed.

BZ#1100865

Support for syslog was missing from the Red Hat Enterprise Virtualization Hypervisor installer which meant syslog was not able to be configured during auto-installation. Now, the syslog parameter is supported as expected and syslog can be configured with the auto-installer using kernel arguments.

BZ#1073724

The 'edit-node --update' command can now be used to update individual packages in a Red Hat Enterprise Virtualization Hypervisor ISO.

BZ#1039231

Support for Broadcom Corporation NetLink BCM57780 Gigabit Ethernet PCIe has been added to the Red Hat Enterprise Virtualization Hypervisor.

BZ#1086268

With the Red Hat Enterprise Virtualization 3.5 release, you can now use a Red Hat Enterprise Virtualization Hypervisor 7.0 in your Red Hat Enterprise Virtualization environment. The Red Hat Enterprise Virtualization Hypervisor 7.0 is a minimal operating system based on Red Hat Enterprise Linux 7.0 that is designed to provide a simple method for setting up a physical machine to act as a hypervisor in a Red Hat Enterprise Virtualization environment. The minimal operating system contains only the packages required for the machine to act as a hypervisor, and features a simple text user interface for configuring the machine and adding it to an environment.

BZ#960379

Partial configuration of kdump is now supported during auto-installation of the Red Hat Enterprise Virtualization Hypervisor. 'kdump_local=1' can now be used to store core dumps locally, and 'kdump_ssh' and 'kdump_ssh_key' can be used to configure kdump for remote storage of the core dump.

BZ#1018063

Previously, a new CIM group added in the %post script created ownership problems of plugins when upgrading from a pre-plugin version of Red Hat Enterprise Virtualization. Now, the %post script logic has been improved to avoid this problem.

BZ#1123413

Previously, kdump was not started on boot and prevented the kdump fencing feature from working on the Red Hat Enterprise Virtualization Hypervisor. Now, kdump is started by default and works with the Hypervisor as expected.

Chapter 4. RHBA-2015:0159 vdsm

The bugs contained in this chapter are addressed by advisory RHBA-2015:0159. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2015-0159.html.

vdsm

BZ#1098769

Previously, the default multipath configurations for EMC-VNX storage were conflicting with those required for VDSM. This caused recovery operations in the event of a failure to be too slow or, in certain cases, to never end. Now, VDSM is shipping with a new multipath configuration that overrides the previous defaults from the multipath package. As a result, recovery operations are now on par with all the other storage solutions.

BZ#1090109

Previously, the start time for a virtual machine's maximum migration time was calculated too early. When more than three migrations were performed, later migrations would fail due to timeouts. This problem has been now resolved.

BZ#1144639

Previously, the VDSM registration service was not updated to include the new VDSM persistence scheme, meaning that VDSM registration would not persist its bridge configuration. VDSM registration now uses the new persistence scheme, so that after registration the management bridge persists as it should.

BZ#1152587

The issue_lip operation has been found to be disruptive on some storage servers, causing storage connection issues. Domains became inaccessible on random occasions. With this update, the issue_lip operation is disabled by default. As a result, discovering new LUNs on Fibre Channel storage server is not supported by default. Users can enable this option through new VDSM configuration (hba_rescan) if this option is compatible with the storage server. A future Red Hat Enterprise Virtualization version will support discovering new LUNs by default.

BZ#1159839

Previously, Red Hat Enterprise Linux Virtualization 3.4.0 introduced a regression where SCSI scans for FC devices were disabled. As a result, new LUNs on FC servers were not discovered automatically. With this release, the SCSI scan for FC devices has been reintroduced so that new LUNs on FC servers can be discovered automatically.

BZ#1162784

With this release, VDSM private keys are no longer collected by sosreports.

BZ#1053114

Previously, extracting information on networks took a long time when there were multiple networks defined on the host. Using a host with 200+ networks was very slow or impossible. Now, the code has been refactored with attention to asymptotic time efficiency, so that 1000 networks are workable.

BZ#1072030

Previously, missing unit conversion caused the reported shared memory amount to be much higher than expected. Proper unit conversion has now been added, resulting in accurate shared memory amount reporting.

BZ#1097674

Previously, a change was introduced that prevented a bond's IP from being cleared when a new one was being set, causing the new IP address to be configured as secondary, alongside the previous one, instead of as the sole primary IP address. Now, the previous IP bonding configuration is removed when reconfiguring the IP, so that there are no leftover IPs from previous configurations.

BZ#1173257

Previously, it took a couple of seconds to collect the memory and balloon information of a virtual machine that had just finished migration. This caused MOM to receive zeros and subsequently try to set the balloon size to zero. The guest operating system then returned all memory it could and crashed with kernel panic once the kernel needed to allocate some buffer. Now, VDSM does not report any ballooning information (not even zero) until it collects the necessary data, so migrating ballooned hosts works properly.

BZ#1136982

Previously, a naming issue caused a permanent communication disruption with certain guest virtual machines until the complete stop and restart of the affected virtual machine. This problem has been now resolved.

BZ#1101021

Previously, a timing issue caused a permanent communication disruption with all guest virtual machines until the restart of the VDSM process. This problem has been now resolved.

BZ#1121295

Previously, the number of sockets configured for a virtual machine exceeded the QEMU limits. This caused the virtual machine to fail to run. Now, VDSM sends a proper socket value number according to configured limits inside the Engine (MaxNumOfVmSockets), so the virtual machine runs with a proper CPU topology that meets the vCPU required.

BZ#1113948

Previously, when a device was removed while sampling network device data, it would raise an Exception, invalidating the whole sampling. Now, when failing to retrieve data for a network device due to it being removed, it is deleted from the output so that the rest of the sampling can continue.

BZ#1111234

With this release, hosts can now keep a connectivity history log. Sometimes, hosts fail due to transient connectivity failures. To help debug these failures, users can now check /var/log/vdsm/connectivity.log. This log can be reviewed via the log inspector. It includes changes in interface operational status, speed, and duplex. It also reports when the Engine is disconnected from the host and stops polling it by logging client_seen:False.

BZ#1104774

With this release, Red Hat Enterprise Virtualization Manager now allows up to 4000 GB of RAM per virtual machine.

BZ#1062617

Previously, VDSM's netinfo.py was reporting speed 0 for the VLAN devices in the system. Now VLAN devices, just like NICs and bonds, will have a speed reported (the same speed as that of their underlying device). Networks that are defined on top of VLANs will now have a non-zero speed associated with them.

BZ#821493

When a multi-processor virtual machine communicates with other virtual machines on the same host, its CPU may generate traffic faster than a single virtio-net queue can consume it. This feature aims to avoid this bottle neck by allowing multiple queues per virtual network interface. Note that this is effective only when the host runs a Red Hat Enterprise Linux 7 kernel >= 3.10.0-9.el7.

BZ#1092166

Previously, it was impossible for a third-party tool to get access to VDSM images. It is now possible to prepare and teardown images

(not in use by a virtual machine) in order to inspect the content.

BZ#1125237

Previously, the logging level of libvirt was set to debug mode, which greatly increased log file size and negatively impacted performance for production environments. Now, the default logging level of libvirt is used and verbosity is decreased. If /run/systemd/journal/socket exists on the machine, libvirt's log file may be changed to journal. Refer to http://libvirt.org/logging.html for more information on the journal change.

BZ#1100264

When a virtual device such as a VLAN was created with a pre-existing associated ifcfg file, Udev would execute ifup on it. If the ifcfg file was configuring the VLAN device for DHCP network address management, when VDSM executed ifup for the device, dhclient would fail as another dhclient (the one started by the ifup Udev performed) would already be controlling the device. Now, hotplug=no has been added to the ifcfg file for virtual devices, so configuring VLANs with DHCP works properly.

BZ#1102549

Previously, VDSM reported an 'ERROR' level log message for the non-severe and common event of a guest virtual machine disconnecting from the communication channel. This can happen when a virtual machine reboots, shuts down, or gets suspended or migrated. Now, the message has been changed to an 'INFO' level log message.

BZ#1113185

Previously, users who wanted to take advantage of specific mount options their storage array supported could only define a POSIX domain, thus losing the enhancements Red Hat Enterprise Virtualization provides for NFS. With this release, it is now possible to specify custom mount options for NFS storage domains.

Chapter 5. RHSA-2015:0158 Moderate: Red Hat Enterprise Virtualization Manager 3.5.0

The bugs contained in this chapter are addressed by advisory RHSA-2015:0158. Further information about this advisory is available at https://rhn.redhat.com/errata/RHSA-2015-0158.html.

ovirt-engine-backend

BZ#1154630

Red Hat Enterprise Linux guests do not support NIC hot plugging by default. Install powerpc-utils version >=1.2.19 on the guest to enable NIC hot plugging.

BZ#1083998

With this update, you can now use Foreman to detect bare metal hosts, allowing the administrator to select and provision the bare metal host as a Red Hat Enterprise Virtualization Manager host.

BZ#1153544

Previously, after a failed migration, the target virtual machine remained in a locked state. Further operations on the virtual machine failed with an error: 'Cannot run VM. VM <VM NAME> is being migrated.' With this update, locked virtual machines are released upon migration failure and further operations on the virtual machine are allowed.

BZ#1149135

Previously, updating a virtual machine from a pool that was set to use the latest version of the template on which the pool is based would sometimes fail. This resulted in virtual machines that could not be updated to the latest version of the template being removed from the pool. Now, the version of the template for virtual machines in pools has been corrected so that virtual machines are no longer removed from pools under these circumstances.

BZ#1148623

Previously, the America/New_York time zone was mapped without taking the daylight savings time into account. This would result in the time of virtual machines shifting by one hour on certain dates. Now, the mapping of the America/New_York time zone takes daylight savings time into account.

BZ#1140430

Previously, failing to attach an ISO domain caused the storage pool manager to fail over and start an attempt to select a new storage pool manager. This behavior caused a potential fail over

storm if the domain itself was corrupted, leaving the system without a storage pool manager for a prolonged time. With this update, a failed attempt to attach an ISO domain to a data center triggers an error message, but does not cause the storage pool manager to fail over.

BZ#1044042

With this feature, users can now configure bridging options from Red Hat Enterprise Virtualization Manager.
Previously, the Manager only configures a small subset of values of a linux bridge. Users who made customized configuration changes would find the configuration overridden by the Manager. Bridging options can now be supplied when provisioning a network on a host using the "bridge_opts" key. These custom properties are accessible through the Administration Portal, REST API, and software development kits.

BZ#1134009

With this update, network labels can now be added to networks that are being used by running virtual machines.

BZ#987295

With this release, support for periodic power management health check to detect and warn about link-down detection of power management LAN has been added.

BZ#977306

This enhancement adds information about password validity to console.vv files. It affects 'Native client' console invocation for SPICE and VNC.

BZ#1044033

With this feature, you can now configure ethtool options from Red Hat Enterprise Virtualization Manager. Previously, the Manager only configures a small subset of the values of a network interface. Users now have the option to use the ethtool utility to customize their usage of network interface. The engine-config tool has to be used initially for the "ethtool_opts" key to be made available. These custom properties are accessible through the Administration Portal, REST API, and software development kits.

BZ#1123754

Upon creating a new DirectLUN disk, the LUN visibility on a host is now validated. If the specified LUN isn't visible to the host, the action would be aborted and a proper error message returned. Note that the validation is only executed if a host is specified by the user; otherwise, no validation is performed.

BZ#1157211

Previously, memory and CPU resources that were reserved for a migrated virtual machine on the destination host were not cleared when a migration failed. With this update, the reserved memory and CPU resources are now cleared properly upon migration failure.

BZ#1120858

This enhancement adds the ability to disable fencing for a cluster. This allows system administrators who are aware that certain hosts in a cluster may experience temporary connection issues to disable and re-enable fencing when performing maintenance on a machine.

BZ#1112359

Previously, virtual machines would be reported as running on the wrong host after failing to migrate due to a maintenance operation on the host. This would prevent hosts where such virtual machines were reported as running from being removed from the Manager. Now, virtual machines are reported as running on the correct host, and it is possible to remove hosts correctly when there are no running virtual machines on those hosts.

BZ#1097256

Previously, virtual machines that failed to migrate to another host due to a maintenance operation on a host would cause deadlocks in the engine database. This would result in maintenance operations taking a long time to complete when virtual machines failed to migrate. Now, deadlocks no longer occur, allowing maintenance operations to complete more quickly when virtual machines fail to migrate.

BZ#1104195

Previously, virtual machines that went down on a destination host as part of a migration operation were considered as having crashed. This would result in an incorrect audit log entry stating "Domain not found: no domain with matching uuid". With this update, The Manager no longer treats virtual machines that went down on a destination host as having crashed, preventing incorrect audit log entries from being recorded when a virtual machine goes down during a migration operation.

BZ#1096971

Previously, when importing an ISO domain or an export domain to a data center, the imported domain was activated right after being attached to the data center. With this update, an imported domain will not be activated by default unless the corresponding check box in the import domain dialog is checked.

BZ#1091692

Previously, when removing an labeled network from a data center directly, or removing an labeled network from a cluster first, then from the data center, the behavior was inconsistent. In the latter scenario, the network became an unmanaged network. With this update, when removing a labeled network, the behavior is the same in both scenarios and would not cause any networks to be left in an unmanaged state.

BZ#999975

Previously if a vlan device had a non-standard name (the standard is- "dev.VLANID"), the engine couldn't handle and display it. This feature adds the functionality to display such vlan devices.

p.s- those devices are just displayed in host->network interfaces sub tab. Setup networks operations cannot be performed on them.

BZ#1043808

Previously, for an host interface that has multiple VLAN interfaces, the highest MTU available was assigned to all VLAN interfaces under that interface, and caused the host going into a non-responsive state. This bug fix moves the setting of the host level value of a default MTU to the engine side so a default value is in place if the MTU is not manually set. You can set the default MTU by setting the 'DefaultMTU' property using the engine-config tool. The default host level MTU must be the same as the data center level MTU, otherwise the network is considered out of synchronization. After upgrading to Red Hat Enterprise Virtrualization 3.5, if the host level and the data center level MTU is not the same, the network will be out of synchronization.

BZ#1133561

Previously, when a user tried to stop a virtual machine which is already down, an error message returned saying the operation cannot be done. With this update, a DEBUG/INFO message is displayed instead of an error message.

BZ#1156577

Previously, upgrading LDAP to rhel-6.6 would cause the Manager to fail to communicate with IPA due to a NegativeArraySizeException exception under certain circumstances. This would prevent LDAP authentication from functioning correctly. As a work around, you can explicitly set "minssf=1" on the IPA side to enable Java to communicate correctly, or run the following command on the machine where the Manager is installed to protect only the authentication sequence:

```
# engine-config -s SASL_QOP auth
```

BZ#1087745

The 32-bit memory limit was used for 64-bit Red Hat Enterprise Linux and Windows operating systems. As a result, a warning was displayed when trying to configure a virtual machine with memory exceeding the 32-bit limit but valid for 64-bit systems. With this update, 64-bit Red Hat Enterprise Linux and Windows operating systems use appropriate limits and no warning is returned for the correct limits.

BZ#1097622

When running a virtual machine with a DirectLUN disk attached using a VirtIO interface, the Manager now sets the 'device' property to 'lun' so that the 'disk' tag resembles [1], allowing generic SCSI commands from the virtual machine to be properly accepted and passed through to the physical device. Previously, the correct attribute was sent only on hot-plug.

[1] <disk type='block' device='lun' snapshot='no'>

BZ#920708

Previously, creating a new storage domain would fail if the given path was to a pre-existing domain. With this update, importing existing domains and adding new domains are separated as two actions and users can now create a new storage domain (NFS) on a mount that has existing storage domains. See the Technical Guide, XML Representation of a Storage Domain for an example. Also see BZ#716511 for more information on this feature.

BZ#1115845

LUN information synchronization [1] is now invoked whenever the status of a storage domain changes to 'Active' (for example, when a storage domain is detected as active upon activating the storage pool manager). Previously, this process was activated only when manually activating a domain.

[1] The process of synchronizing LUN information from the underlying storage with the engine database, such as when adding, removing, or extending a LUN in storage, is properly reflected in the engine database and consequently in the user interface and REST API.

BZ#890517

With this update, glusterVolumeProfileInfo is now supported as part of the Gluster profile support.

BZ#922377

Previously, only certain virtual machine properties could be updated while the virtual machine was running. Other properties

could only be updated while the virtual machine was down. Now, all properties can be updated on a running virtual machine, but those that cannot be applied immediately will be saved and applied the next time the virtual machine is shut down.

BZ#1118847

Previously, the Red Hat Enterprise Virtualization Manager was configured to set all virtio-SCSI direct LUN devices to the "LUN" device type. This device type does not support direct LUN read-only capability. Now, the Manager sets virtio-SCSI direct LUNs to the "disk" device type when the read-only option is enabled, which enables read-only functionality via SCSI emulation. This functionality is important, in particular, for Cloud Forms Management Engine appliances attempting to run smart-state analysis against Red Hat Enterprise Virtualization data storage domains with a large number of backing LUNs.

BZ#1093393

This release introduces a change to the iSCSI multipath bond to block the addition of required networks to the bond. In previous releases, required networks were allowed to be added to the iSCSI multipath bond, and could cause a host to become non-operational even if one of the networks were lost.

BZ#1025376

The 'Change CD' window now displays the name of the CD that is currently attached to a virtual machine.

BZ#1129634

Previously, sparse (thinly provisioned) virtual machine disks that were imported from a file storage domain to a block domain would change format to COW preallocated. Disk images in this format could not be exported, because the disk configuration was incompatible with the storage domain type. Now, a fix introduced in https://bugzilla.redhat.com/show_bug.cgi?id=1116486 converts the images to COW sparse instead, and images can be successfully exported.

BZ#1092884

Previously, virtual machine migration time was displayed in seconds, even for large values. Now, migration time is displayed in hours, minutes, and seconds.

BZ#1119922

A new option in the 'Fencing Policy' tab of the 'New/Edit Cluster' window allows users to disable fencing for any host that has storage connectivity. This is useful to prevent fencing in cases where a host that uses storage has a network issue, but the

services it provides may still be available.

BZ#947965

Previously it was possible to remove virtual machines during powering down or migrating states. Now, that action is properly configured and users can remove virtual machines only when the machines are in "Down" state.

BZ#1114253

Previously, a host performing a fencing operation had to be in the same data center as the host being fenced. Now, a host can be fenced by a host from a different data center.

BZ#1120829

A new option in the 'Fencing Policy' tab of the 'New/Edit Cluster' window allows users to disable fencing of hosts in the cluster if more than a user-defined percentage of hosts have connectivity issues. This can prevent hosts being fenced in scenarios where hosts are in a 'Non-Responding' or 'Connecting' state due to a general network connectivity error, rather than a host error.

ovirt-engine-config

BZ#1128949

Two new configuration values for how OVFs are stored on storage domains were added and exposed to users in the engine-config tool: 'OvfUpdateIntervalInMinutes', which controls how often (in minutes) the update process is run; and 'OvfItemsCountPerUpdate', which controls how many virtual machines OVFs are saved per single VDSM call.

ovirt-engine-restapi

BZ#1062435

With this update, users can now add, update, and delete scheduling policies through the REST API.

BZ#1076705

Previously, the Red Hat Enterprise Virtualization Manager Command Line Interface (rhevm-shell), which is built on top of the Python SDK, did not support assigning custom scheduling policies to a cluster. The same limitation existed in the REST API. Now, custom scheduling policies can be assigned using a reference, by name or by ID, to the new '/schedulingpolicies' collection. The new options are '--scheduling_policy-name' and '--scheduling_policy-id'. The example below assigns a custom policy (by ID) to

```
'mycluster':

# update cluster mycluster --scheduling_policy-id 00000000-0000-
0000-0000-000000000000
```

BZ#1093784

Previously, requests sent via the REST API containing the Expect header did not have the intended effect. This meant that requests that used the Expect header to indicate that they require synchronous execution were actually executed in an asynchronous fashion. This behavior was expected, because the Apache web server rejects a request with an Expect header if it contains any value other than '100-continue'; to mitigate this, the Red Hat Enterprise Virtualization Manager explicitly removed the header from every request. Now, the Manager has been modified to accept an alternative X-Ovirt-Expect header, which has the same values and semantics as the Expect header. To ensure that this header has the desired effect, users must send both the Expect and the X-Ovirt-Expect header with the same value. Developers of client software are encouraged to modify their applications to send both headers with the same value, so that requests will work with previous and upcoming versions of the Manager.

BZ#1101565

Hosts can now be approved via the REST API, as well as via the UI.

BZ#1101018

Custom preview snapshot is now supported in the REST API. Optional 'restore_memory' and 'disks' tags are now accepted in the 'preview_snapshot' action.

BZ#1103490

An exception was thrown when the virtual machine statistics was accessed using the REST API, and it was not possible to retrieve the statistics. With this update, users can now access the statistics using the REST API.

BZ#996512

Users can now log in to a virtual machine (with guest agent installed) via the REST API, using the new 'logon' action. This functionality was already available in the UI. The Manager sends the login credentials to the guest agent, which starts a session of the guest operating system and unlocks the display.

ovirt-engine-setup

BZ#1103676

Previously, the Red Hat Enterprise Virtualization Manager stored files in /var/tmp/ovirt-engine, and tmpwatch deleted those files on a monthly basis. This caused a failure in the Manager. Now, files are stored in /var/lib and tmpwatch does not delete them.

BZ#988422

Red hat Enterprise Virtualization has incorporated the OpenStack Neutron service as a network provider as part of the 3.4 release. However, to provision Neutron services, users need to manually deploy Neutron and Keystone services. With this update, users can now download the Neutron Appliance to deploy a Red Hat Enterprise Linux 7.0 based virtual machine with Neutron installed. The Neutron Appliance was designed to simplify the deployment process.

BZ#1109326

During upgrades, if automatic firewall configuration with iptables was chosen, NFS server ports were closed off. This caused problems for NFS storage domains. Now, NFS status is checked before iptables configuration is generated.

BZ#985945

The Red Hat Enterprise Virtualization Manager websocket proxy can now be installed and configured (via engine-setup) on a separate machine from the machine on which the Manager is installed.

BZ#1125834

Previously, when the specified storage path for an ISO domain had UUID validation exceptions (i.e. was not empty), an unclear error message was shown:
[ERROR] Cannot access mount point /exports/iso/: badly formed hexadecimal UUID string
Now, a more explanatory error is given:
[ERROR] Cannot access mount point /exports/iso: Error: directory /exports/iso is not empty

BZ#1103976

Previously, setting up a PostgreSQL database with the engine-setup command generated weak passwords for PostgreSQL users. Since the PostgreSQL database is accessible remotely with a default Red Hat Enterprise Virtualization Manager installation, this was a security issue. With this update, stronger random passwords were generated and the password length has been extended to 22 characters.

ovirt-engine-userportal

BZ#1001419

In the User Portal, the right-hand pane took up too much space

and left less room for virtual machines. When users resized the right-hand pane, a scrollbar was generated and hid the "Edit" button for "Console" options. With this update,the right-hand pane was restructured so a horizontal scrollbar is not produced when the right-hand pane is resized smaller.

BZ#955235

With this feature, BIOS boot menu for virtual machines is now supported. This feature eases selecting boot options when needed.

BZ#1085380

Previously, when trying to save a virtual machine editing window in the User Portal with validation errors, users cannot see the error message and would try to save it with a failed attempt. The error message was only available if the user clicked on 'Show Advanced Options'. With this update, users can now see the advanced options and the error message when attempting to save with validation errors.

ovirt-engine-webadmin-portal

BZ#1116486

Since Red Hat Enterprise Virtualization 3.4, mixed domain types in the same data center were allowed. To be able to move or copy a image from a file domain to a block domain, the images were converted to raw+preallocated in the operation and was resource wasteful. With this update, images are converted to cow+sparse instead to be more resource efficient. This feature is also backported to Red Hat Enterprise Virtualization 3.4.2.

BZ#1064273

Previously, when changing the data center for a host to a different data center with existing hosts, an attempt to create a virtual machine in the Administration Portal would fail. With this update, users can now create a virtual machine under the mentioned circumstance.

BZ#1085136

With this release, a disk's description property can be changed while the virtual machine is running. A description of a disk may change frequently (for example, when you install new software on the guest), and having to shut the virtual machine down in order to update it can hinder production needs.

BZ#1043430

With this update, Firefox 31 is added as a supported browser.

BZ#1064544

With this release, the graphical user interface for the Administration Portal and User Portal has been updated to provide Red Hat customers with better unified interface experience across products. After upgrading to Red Hat Enterprise Virtualization 3.5, clear your browser cache to see the updated interface.

BZ#1028387

Previously, virtio-serial and balloon devices were treated as unmanaged devices on Windows virtual machines. As a result, the addresses of the virtio-serial and balloon devices changed every time a virtual machine was started and users were asked to install drivers again. With this update, virtio-serial and balloon devices are now managed devices. Users are no longer asked to install drivers for the virtio-serial and balloon devices every time a Windows virtual machine is started.

BZ#1114241

With this update, when editing the "Setup Host Networks" window, the "Save Network Configuration" check box is now marked by default to prevent user configuration changes wiped out by accident.

BZ#1098638

Previously, if smartcard support was enabled on a template, every time when the template was edited and saved, a new smartcard entry was created. This eventually caused virtual machines to fail to boot. With this update, only one smartcard device is available for templates that have smartcard support enabled.

BZ#1048019

This feature optimizes queries for data associated with the system tree. Previously, the queries for data were serialized, so one would not start before the previous one was completed, even though there was no relationship between them. Now the queries run in parallel, improving UI start up time.

BZ#1092609

With this feature, users can now search for objects that have tags or objects that do not have tags.

BZ#859024

When performing actions such as unplugging a Virtual NIC, a confirmation dialog is displayed to prevent user performing the action by accident.

BZ#804530

This feature changes the "Slot" field to "Service Profile" when cisco_ucs is selected as the fencing type.

BZ#1053884

Previously, a user-paused virtual machine was unable to be migrated. Now, this issue has been fixed and user-paused virtual machines can migrate.

BZ#1131693

The fix allows Network Level Authentication to be used with Native Remote Desktop Protocol (RDP) client. Note that Network Level Authentication is still disabled for RDP browser plug-in.

BZ#1123396

Previously, infrastructural GUID computation for certain entities was highly inefficient. When many virtual machines had to be displayed in the specified sub-tab, this inefficient computation became visible as the browser would wait on it to display the virtual machines. This caused general sluggishness in the browser, and sometimes triggered an "unresponsive script" error message. Now, the GUID computation has been optimized so that the tab data is loaded as fast as other tabs with comparable data sets (the Virtual Machines main tab, for example).

BZ#1121454

Previously, creating or editing an NFS storage domain's mount path so that the server's name ends in a digit (for example: myhost1:/path/to/data) was not allowed, meaning that legal host names could not be used as storage servers providing NFS storage to Red Hat Enterprise Virtualization. This limitation has now been removed so that hosts with names ending in digits can now be used as NFS servers.

BZ#1070823

With this feature, you can now edit the "Wipe after Delete" property of a disk even while the virtual machine is running.

BZ#1013670

Previously, the comment field was not set up properly for creating a template from a virtual machine so comments were not saved. With this update, the comment field is properly saved when creating a template from a virtual machine.

BZ#1100194

Previously, when using an Internet Explorer 9 browser, it was not possible for a user to use a mouse to select a specific template if there were more than three templates available. With this update, users can now use a mouse to scroll down the list of templates available.

RFEs

BZ#906938

With this update, support for storage quality of service has been added.

BZ#906927

With this update, support for CPU quality of service has been added.

BZ#987299

With this update, you can now set event notifications for NIC slave or bond faults, provided there is a network or label on the interface. Four new events have been made available for selection to configure your event notifier. They are: HOST_INTERFACE_STATE_UP, HOST_INTERFACE_STATE_DOWN, HOST_BOND_SLAVE_STATE_UP, and HOST_BOND_SLAVE_STATE_DOWN. To enable or update your event notifier, subscribe to ovirt-engine-notifier to receive notifications on your selected events. See the Administration Guide, Configuring Event Notifications for more information.

BZ#874328

With this enhancement, a new instance management screen is now available in the Administration Portal.

BZ#828591

Administrators can now identify the optimal balance of virtual machines within a cluster. In addition, administrators can determine how to place new virtual machine workloads into a cluster with enough total available resources, and avoid scenarios whereby no single host has enough resources for a new virtual machine.

BZ#817180

With this release, MachineObjectOU is now available for configuration for virtual machines that are using Sysprep. This allows users to specify an Active Directory OU for virtual machines to join.

BZ#1110636

With this update, you can now install a RHEL-based Hypervisor on IBM Power 8 hardware for your Red Hat Enterprise Virtualization environment.

BZ#962880

With this enhancement, when a grid is loaded and it contains only one item and it can be selected, the system will automatically select that item and display the detailed information of that item. This feature saves the user from manually clicking the item.

BZ#723211

With this feature, users can now clone a virtual machine directly from an existing virtual machine without the need to create a template first making the process more time and resource efficient.

BZ#895222

In the Administration Portal, users can sort tables by clicking on column headers.

BZ#716511

Red Hat Enterprise Virtualization 3.5 provides support for migrating storage domains amongst different Red Hat Enterprise Virtualization data centers or different deployments. This functionality allows the transfer of virtual machines between setups without the need to copy the data into and out of an export domain, or the need to recover after the loss of an engine database. Also see BZ#920708 for the REST API implementation.

BZ#1102018

Previously, the OpenStack Networking (Neutron) integration supported both the Linux Bridge and Open vSwitch plug-ins. Since Open vSwitch is the recommended plug-in to use with Red Hat Enterprise Linux OpenStack Platform and it offers feature parity with Linux Bridge, the Linux Bridge plug-in is dropped from the integration.

BZ#918138

With this enhancement, it is now possible to configure serial numbers for virtual machines on three different levels: engine-config level, cluster level, and virtual machine level. At each level, three modes of serial numbers are available: use host UUID (legacy), use VM UUID, and provide custom serial number _____.

BZ#988392

With this update, users now have the option to dismiss unwanted alerts from the Administration Portal.

BZ#906243

This features adds the ability to configure the host name of a virtual machine using sysprep.

BZ#962220

This feature adds the ability to configure the system, user, and machine locale for a virtual machine using sysprep in the New Virtual Machine and Edit Virtual machine window.

BZ#1034885

With this update, users can now see the overview of snapshots in the Administration Portal.

BZ#800155

This features adds the ability to disable copying and pasting to virtual machines through SPICE connections, allowing administrators to restrict this functionality due to security reasons. This functionality is enabled by default.

BZ#967466

In the Administration Portal, a progress bar is now available to indicate the progress of migrating a running virtual machine.

BZ#1022795

Previously, when creating a virtual machine disk , the Administration Portal suggested a default disk alias consisting of 'VMname+_disk+number. The number came from the total amount of virtual machines disks plus 1. For example, for a virtual machine name 'V1' with 3 existing disks, the suggested name was'V1_disk4. However, the suggestion mode didn't recycle disks aliases correctly. If a disk was deleted, the number was not reused. With this update, the suggestion mode recycled the unused numbers to form new virtual machine disk aliases.

BZ#1016916

With this update, you can now search for virtual machines in the Administration Portal using their MAC addresses.

BZ#1059435

This feature allows administrators to take full advantage of Self-Hosted Engine and can implement Self-Hosted Engine on a Red Hat Enterprise Virtualization Hypervisor Host.

BZ#1015186

When a block storage domain exceeds a certain number of logical volumes defined in a configuration value, each action that results in the creation of a new logical volume on the domain will add an audit log warning that the number of logical volumes on that domain has exceeded the defined number. The number of logical volumes is defined in the configuration value 'AlertOnNumberOfLVs' and its default value is 300.

BZ#1077284

With this feature, users can now configure a MAC pool with larger address ranges. Red Hat recommends to configure the MAC address pool to contain the majority of MAC addresses to be used. Only MAC addresses defined in the MAC address pool will be stored in memory efficiently.

BZ#879077

Previously,when an entity changed, the system tree was not automatically refresh and users had to press the refresh button to update the tree. With this update, when an entity changes, the system tree will refresh automatically.

BZ#878662

With this update, you can now set up custom fence agents for your Red Hat Enterprise Virtualization 3.5 environment. For more information, see https://access.redhat.com/articles/1238743.

BZ#1067162

The Hosted Engine can now be deployed on iSCSI storage domains.

BZ#1065753

With this update, users are asked to optionally specify a reason when performing maintenance operations on a virtual machine. The feature can be set in the cluster properties to make the function optional or not.

BZ#1113937

With this update, the engine can now integrate with Apache authentication, for example mod_auth_kerb, to accept users already authenticated by Apache and enable single sign-on to the User and Administration Portals.
Note that this feature conflicts with the password delegation

feature in 3.4 (also known as the single sign-on to virtual machine feature) as the engine does not have access to user passwords anymore, the password cannot be delegated to virtual machines.
Also note that when this feature is used, the sign out button in the User Portal and Administration Portal will not work. The user will remain logged in even after clicking the sign out button. To sign out, the user must sign out from the single sign-on provider. For more information on configuring this feature, see the ovirt-engine-extension-aaa-ldap package documentation[1].

[1] http://gerrit.ovirt.org/gitweb?p=ovirt-engine-extension-aaa-ldap.git;a=blob;f=README;hb=HEAD#l141

BZ#1010079

Feature:Virtual NUMA Node

Reason: To provision larger guests

Result (if any): NA

BZ#1010059

Feature: NUMA

Reason: For customers who provision larger guests

Result (if any): NA

BZ#1052348

In order to provide additional tools for debugging and troubleshooting potential storage problems from the Red Hat Enterprise Virtualization Hypervisor, the iotop package is now included in Red Hat Enterprise Virtualization Hypervisor images.

BZ#1025831

With this update, you can configure the administrator password and organization name in the Initial Run tab of the Run Once menu.

BZ#894084

With this enhancement, a warning message is displayed in the user interface if SELinux is disabled to remind users of the SELinux status.

BZ#977079

This feature adds support for enabling a paravirtualized random number generator (RNG) in virtual machines. To use this feature, the random number generator source must be set at cluster level to ensure all hosts support and report desired RNG device sources.

This feature is supported in Red Hat Enterprise Linux hosts of version 6.6 and higher.

BZ#1038632

This enhancement adds a button to the SPICE-HTML5 page and allows users to display console debug information when needed.

BZ#1032686

Previously, all virtual machine OVFs were stored on the master domain and are being updated asynchronously by the OvfAutoUpdater. With this feature, OVFs are now stored on all wanted domains to provide better recovery ability, and to reduce the use of master_fs and the master domain.

BZ#1083760

With this feature, a host is prevented from rebooting when the host is in the middle of a Kdump process to prevent any log loss.

BZ#1110172

Previously, during fencing, a host cannot be accessed through the network and administrators could not check if the host was working or not. With this update, the host can be checked using the Sanlock lease information. Now, hosts can be checked even during fencing.

vulnerability

BZ#1081896

It was found that the oVirt web admin interface did not include the HttpOnly flag when setting session IDs with the Set-Cookie header. This flaw could make it is easier for a remote attacker to hijack an oVirt web admin session by leveraging a cross-site scripting (XSS) vulnerability.

BZ#1081849

A Cross-Site Request Forgery (CSRF) flaw was found in the REST API. A remote attacker could provide a specially crafted web page that, when visited by a user with a valid REST API session, would allow the attacker to trigger calls to the oVirt REST API.

Revision History

Revision 3.5-2	**Mon 27 Apr 2015**	**Red Hat Enterprise Virtualization Documentation Team**

Updated for the Red Hat Enterprise Virtualization 3.5.1 release.

Revision 3.5-1	**Fri 07 Nov 2014**	**Julie Wu**

Initial creation for the Red Hat Enterprise Virtualization 3.5 release.